

Distinct values of bilinear forms on algebraic curves

Claudiu Valculescu

Frank de Zeeuw

Abstract

Let $B_M : \mathbb{C}^2 \times \mathbb{C}^2 \rightarrow \mathbb{C}$ be a bilinear form $B_M(p, q) = p^T M q$, with an invertible matrix $M \in \mathbb{C}^{2 \times 2}$. We prove that any finite set S contained in an irreducible algebraic curve C of degree d in \mathbb{C}^2 determines $\Omega_d(|S|^{4/3})$ distinct values of B_M , unless C has an exceptional form. This strengthens a result of Charalambides [1] in several ways.

The proof is based on that of Pach and De Zeeuw [8], who proved a similar statement for the Euclidean distance function in \mathbb{R}^2 . Our main motivation for this paper is that for bilinear forms, this approach becomes more natural, and should better lend itself to understanding and generalization.

1 Introduction

Pach and De Zeeuw [8] proved that a finite set S on an irreducible algebraic curve of degree d in \mathbb{R}^2 determines $\Omega_d(|S|^{4/3})$ distinct Euclidean distances, unless that curve is a line or a circle. In this paper we prove an analogous result for functions $\mathbb{C}^2 \times \mathbb{C}^2 \rightarrow \mathbb{C}$ of the following form, with $p = (p_x, p_y), q = (q_x, q_y) \in \mathbb{C}^2$:

$$c_1 p_x q_x + c_2 p_x q_y + c_3 p_y q_x + c_4 p_y q_y.$$

We refer to such functions as *bilinear forms*, and write them more compactly as

$$B_M(p, q) := p^T M q$$

with a matrix $M \in \mathbb{C}^{2 \times 2}$. We assume throughout that M is invertible. For $S \subset \mathbb{C}^2$, we write $\mathcal{B}_M(S) := \{B_M(p, q) : p, q \in S\}$, so $|\mathcal{B}_M(S)|$ is the number of distinct values of B_M on S .

Two particular functions that we are interested in are B_I and B_A for

$$I := \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad A := \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Over \mathbb{R} , $B_I(p, q) = p^T q$ is the dot product, and $B_A(p, q)$ is twice the signed area of the triangle spanned by p , q , and the origin. Distinct values of the dot product on various sets were considered in [12] and [3, Chapter 9], but have not been considered on algebraic curves before. For triangle areas, Charalambides [1] proved (among other results) that for S contained in an algebraic curve of degree d in \mathbb{R}^2 , one has $|\mathcal{B}_A(S)| = \Omega_d(|S|^{5/4})$, unless the curve is a line, an ellipse centered at the origin, or a hyperbola centered at the origin. We improve Charalambides's bound to $\Omega_d(|S|^{4/3})$, give an explicit dependence on d , and extend our bound to general bilinear forms as well as to curves in \mathbb{C}^2 .

The class of curves for which our bound does not hold is actually somewhat larger than for Charalambides, so, strictly speaking, we do not quite improve his bound in all cases. But we show that our class of exceptional curves is best possible for general bilinear forms. This class is captured in the following definition.

Definition 1.1. We call an algebraic curve in \mathbb{C}^2 a *special curve* if it is a line, or it is linearly equivalent to a curve defined by an equation of the form

$$x^k = y^\ell, \quad \text{with } k, \ell \in \mathbb{Z} \setminus \{0\}, \quad \gcd(k, \ell) = 1.$$

We say that two curves C, C' are *linearly equivalent* if there is an invertible matrix $D \in \mathbb{C}^{2 \times 2}$ such that $C' = DC := \{Dp : p \in C\}$. Because k and ℓ are assumed to be coprime, all special curves are irreducible. When k or ℓ is negative, one obtains a more natural polynomial equation after multiplying by an appropriate monomial. Thus special curves include hyperbola-like curves of the form $x^k y^\ell = 1$ with coprime $k, \ell \geq 1$. Ellipses centered at the origin are also included, since these are linearly equivalent to the unit circle $(x - iy)(x + iy) = 1$, which is linearly equivalent to $xy = 1$. Thus all the exceptional curves of Charalambides are special.

We now show that for any special curve, there is a bilinear form that takes only a linear number of distinct values on it.

Example 1.2. If C is special, there are $M \in \mathbb{C}^{2 \times 2}$ and $S \subset C$ such that $|B_M(S)| = O(|S|)$.

- Let C be a line $y = c$. Then $S = \{(2^i, c) : i = 1, \dots, |S|\}$ has $|\mathcal{B}_I(S)| = O(|S|)$.
- Consider the curve C given by $x^k = y^\ell$. Take

$$S := \{(2^{\ell i}, 2^{ki}) : i = 1, \dots, |S|\} \subset C.$$

Then $B_I((2^{\ell i}, 2^{ki}), (2^{\ell j}, 2^{kj})) = (2^\ell)^{i+j} + (2^k)^{i+j}$, so $|\mathcal{B}_I(S)| = O(|S|)$.

- For any other special curve C' , there is an invertible matrix D such that $C' = DC$, for a curve C defined by $x^k = y^\ell$ or $y = c$. Then we can choose $S \subset C$ as above, so that for $p, q \in S$, we have

$$B_M(Dp, Dq) = p^T D^T M D q.$$

Choosing $S' = DS \subset C'$ and $M = D^{-T} D^{-1}$, we have $|\mathcal{B}_M(S')| = |\mathcal{B}_I(S)| = O(|S'|)$.

Our main theorem says that these special curves are the only curves on which B_M could have a linear number of distinct values, while on any other curve B_M must take significantly more values. See Section 4 for a discussion of extensions and generalizations.

Theorem 1.3. Let C be an irreducible algebraic curve in \mathbb{C}^2 of degree d , $S \subset C$ a finite set, and B_M a bilinear form as above. If C is not special, then

$$|\mathcal{B}_M(S)| = \Omega(d^{-14/3} |S|^{4/3}).$$

Proof. We outline how the proof is distributed over the paper. By Corollary 2.2 in Section 2, the bound holds if M is invertible and C has $O(d^2)$ automorphisms. By Theorem 3.1 in Section 3, the only curves that do not have $O(d^2)$ automorphisms are the special curves. \square

For clarity we have chosen not to state our result in the most general form possible. The proof in fact gives a “bipartite” statement (see Theorem 2.1 and also [8]), and can be extended to bilinear functions with linear terms, as well as to reducible curves. We also note that for sets on curves in \mathbb{R}^2 , our proof gives, with a little extra work, a better dependence on d , namely d^{-2} instead of $d^{-14/3}$.

Our proof follows the setup in [8], which is based on that of [10]. It turns out that, for bilinear forms, this setup leads to a more natural and streamlined proof than for the Euclidean distance function in [8]. This was our main motivation for working out this variant in detail, and we hope that it helps to clarify the proof of [8], and increases the potential for generalization. We also wanted to test the limits of this approach, by extending it to complex curves and by explicitly determining the dependence on the degree of the curve. In future work we hope to study more general polynomial functions, as well as functions on curves in higher dimensions.

Let us quickly give the relevant definitions. A set $C \subset \mathbb{C}^2$ is an *algebraic curve* if there is an $f \in \mathbb{C}[x, y] \setminus \{0\}$ such that $C = Z(f) := \{(x, y) \in \mathbb{C}^2 : f(x, y) = 0\}$. The *degree* of C is the minimum degree of a polynomial f such that $C = Z(f)$. The curve C is *irreducible* if there is an irreducible f such that $C = Z(f)$. We frequently use *Bézout's inequality*, which states that the number of intersection points of two distinct irreducible algebraic curves in \mathbb{C}^2 is at most the product of their degrees. In our proof, we also consider algebraic curves in \mathbb{C}^4 ; for their definition, we refer to [4]. A crucial role in the proof is played by linear automorphisms of curves. A *linear automorphism* of an algebraic curve C is an invertible linear transformation $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ such that $T(C) = C$. We often drop the word “linear”.

2 Proof of Theorem 1.3

In this section we give one side of the proof of Theorem 1.3; the other side follows in Section 3. We prove Theorem 2.1, a variant of Theorem 1.3 that is more convenient for the proof, and deduce Corollary 2.2, which, together with Theorem 3.1, directly implies Theorem 1.3.

2.1 A variant of Theorem 1.3

Theorem 2.1 differs from Theorem 1.3 in the following ways. It focuses on the matrix I (i.e., $B_I(p, q) = p^T q$ is the “dot product”), but the statement is slightly more general, in that it bounds the values of the function in a useful “bipartite” way; for $S_1, S_2 \subset \mathbb{C}^2$, it bounds the size of $\mathcal{B}_I(S_1, S_2) := \{B_I(p, q) : p \in S_1, q \in S_2\}$. This more general form allows us to deduce the result for B_M . Finally, the exceptional curves in Theorem 2.1 are those curves that have many automorphisms. In Section 3, we show that the only curves with many automorphisms are the special curves of Definition 1.1.

Theorem 2.1. *Let C_1 and C_2 be irreducible algebraic curves in \mathbb{C}^2 , both of degree at most d , and let $S_1 \subset C_1, S_2 \subset C_2$ be disjoint finite sets with $|S_1| = |S_2| = n$. If C_1 and C_2 each have $O(d^2)$ automorphisms, then*

$$|\mathcal{B}_I(S_1, S_2)| = \Omega(d^{-14/3} n^{4/3}).$$

We first deduce from this theorem a statement that is closer to Theorem 1.3.

Corollary 2.2. *Let C be an irreducible algebraic curve in \mathbb{C}^2 of degree d , $S \subset C$ a finite set, and B_M a bilinear form. If M is invertible and C has $O(d^2)$ automorphisms, then*

$$|\mathcal{B}_M(S)| = \Omega(d^{-14/3} |S|^{4/3}).$$

Proof. We arbitrarily split S into two disjoint sets S_1, S'_2 of the same size (discarding one point if $|S|$ is odd). Then we set $S_2 := MS'_2$. For $p \in S_1, q' \in S'_2$ we have $B_M(p, q') = B_I(p, Mq') = B_I(p, q)$ with $q \in S_2$. We set $C_1 := C$ and $C_2 := MC$. Applying Theorem 2.1 to $S_1 \subset C_1$ and $S_2 \subset C_2$ gives

$$|\mathcal{B}_M(S)| = \Omega(|\mathcal{B}_M(S_1, S'_2)|) = \Omega(|\mathcal{B}_I(S_1, S_2)|) = \Omega(d^{-14/3}|S|^{4/3}). \quad \square$$

2.2 Preparation

In the rest of Section 2 we prove Theorem 2.1. We assume throughout that C_1 and C_2 have $O(d^2)$ automorphisms, so in particular they are not lines.

The matrices in the following definition play an important role in the proof.

Definition 2.3. *Given two points $p_i = (x_i, y_i), p_k = (x_k, y_k) \in \mathbb{C}^2$, we define the matrix*

$$N_{ik} := \begin{pmatrix} x_i & y_i \\ x_k & y_k \end{pmatrix}.$$

To ensure that these matrices behave nicely, we prepare the sets S_1, S_2 as follows.

Lemma 2.4. *There is $S^* \subset S_1$ with $|S^*| \geq n/d$ such that any line through the origin contains at most one point of S^* . Consequently, for any distinct $p_i, p_k \in S^*$ the matrix N_{ik} is nonsingular. Furthermore, there is $T^* \subset S_2$ with the same property and $|T^*| = |S^*|$.*

Proof. For any line L through the origin that intersects S_1 , arbitrarily choose one point of $L \cap S_1$ and remove any other point. Call the result S^* . Since C_1 is not a line, by Bézout it contains at most d points on such a line L , so $|S^*| \geq n/d$.

Similarly pick T^* from S_2 , and remove points from the larger set until $|S^*| = |T^*|$. \square

Notation: The rest of the proof considers only $M = I$, so we write $B := B_I$. We only use the points in S^* and T^* ; we set $m := |S^*| = |T^*|$ and $\mathcal{B} = \mathcal{B}_I(S^*, T^*)$. Throughout this section we denote points of C_1 with the letter p , and points of C_2 with the letter q ; for points of S^* or T^* we similarly use either p_i, p_j, \dots or q_s, q_t, \dots . As said, we assume throughout that neither C_1 nor C_2 is a line.

2.3 Quadruples and curves

To prove the theorem, we find lower and upper bounds on the number of quadruples in

$$\mathcal{Q} := \{(p_i, p_j, q_s, q_t) : p_i, p_j \in S^*, q_s, q_t \in T^*, B(p_i, q_s) = B(p_j, q_t)\}.$$

The lower bound is easily obtained using the Cauchy-Schwarz inequality.

Lemma 2.5. *For \mathcal{B} and \mathcal{Q} as above we have $|\mathcal{Q}| \geq m^4/|\mathcal{B}|$.*

Proof. Write $B^{-1}(b) := \{(p_i, q_s) \in S^* \times T^* : B(p_i, q_s) = b\}$ for $b \in \mathcal{B}$. Then

$$|\mathcal{Q}| \geq \sum_{b \in \mathcal{B}} |B^{-1}(b)|^2 \geq \frac{1}{|\mathcal{B}|} \left(\sum_{b \in \mathcal{B}} |B^{-1}(b)| \right)^2 = \frac{m^4}{|\mathcal{B}|}. \quad \square$$

To obtain an upper bound on $|\mathcal{Q}|$, we relate it to an incidence problem for points and curves in \mathbb{C}^4 . We define algebraic curves C_{ij} and \tilde{C}_{st} in \mathbb{C}^4 as follows: For each pair of points $p_i, p_j \in S^*$, we set

$$C_{ij} := \{(q, q') \in \mathbb{C}^4 : q, q' \in C_2, B(p_i, q) = B(p_j, q')\},$$

and for each pair of points $q_s, q_t \in T^*$, we set

$$\tilde{C}_{st} := \{(p, p') \in \mathbb{C}^4 : p, p' \in C_1, B(p, q_s) = B(p', q_t)\}.$$

Lemma 2.6. *The sets C_{ij} and \tilde{C}_{st} are algebraic curves in \mathbb{C}^4 of degree at most d^2 .*

Proof. The set C_{ij} is the intersection of the irreducible surface $C_2 \times C_2$ and the hyperplane H_{ij} defined by the equation $B(p_i, q) = B(p_j, q')$. This hyperplane does not contain the surface, since then fixing q' would give that C_2 is a line, which we assumed it is not. By [4, Proposition 7.1], it follows that the intersection is one-dimensional, i.e. it is an algebraic curve. By a higher-dimensional affine version of Bézout's inequality (see [4, Theorem 7.7] or [5, Theorem 1]), the degree of this curve is at most $\deg(C_2)^2 \cdot \deg(H_{ij}) = d^2$.

The same arguments apply to \tilde{C}_{st} . □

We have $(q_s, q_t) \in C_{ij}$ if and only if $(p_i, p_j) \in \tilde{C}_{st}$. This suggests that we can think of the curve \tilde{C}_{st} as “dual” to the point (q_s, q_t) , and of (p_i, p_j) as dual to C_{ij} .

Define a point set and a curve set by

$$\mathcal{P} := T^* \times T^*, \quad \mathcal{C} := \{C_{ij} : (p_i, p_j) \in S^* \times S^*\}.$$

Then a point $(q_s, q_t) \in \mathcal{P}$ lies on $C_{ij} \in \mathcal{C}$ if and only if $(p_i, p_j, q_s, q_t) \in \mathcal{Q}$. Thus

$$|\mathcal{Q}| = I(\mathcal{P}, \mathcal{C}) := |\{(p, C) \in \mathcal{P} \times \mathcal{C} : p \in C\}|.$$

It is possible that some C_{ij} coincide as sets, but then we consider them as separate objects.

2.4 Intersections

We want to apply an incidence bound to the points \mathcal{P} and curves \mathcal{C} , and for that we need to control the sizes of the intersections between curves. We define

$$\mathcal{C}_0 := \{C_{ij} \in \mathcal{C} : \text{there is a } C_{kl} \in \mathcal{C} \text{ such that } |C_{ij} \cap C_{kl}| = \infty\}$$

and $\mathcal{C}_1 := \mathcal{C} \setminus \mathcal{C}_0$. Dually, we set

$$\mathcal{P}_0 := \{(q_s, q_t) \in \mathcal{P} : \text{there is a } (q_u, q_v) \in \mathcal{P} \text{ such that } |\tilde{C}_{st} \cap \tilde{C}_{uv}| = \infty\}$$

and $\mathcal{P}_1 := \mathcal{P} \setminus \mathcal{P}_0$. Thus, the curves in \mathcal{C}_0 are “bad” curves that have large intersection with some other curve, while the points in \mathcal{P}_0 are “bad” in a dual sense. We show that the sets \mathcal{C}_0 and \mathcal{P}_0 are relatively small. For the “good” sets \mathcal{P}_1 and \mathcal{C}_1 , the intersections are well-behaved, allowing us to apply an incidence bound.

With these definitions, two fortunate things happen. Whenever curves C_{ij} coincide as sets, they must lie in \mathcal{C}_0 . The curves C_{ii} for any i , which would cause trouble in some of the statements, are also in \mathcal{C}_0 , because they all contain the line $\{(q, q) : q \in C_2\}$. The analogous statements hold for the dual curves and the corresponding points in \mathcal{P} .

We now show that for \mathcal{P}_1 and \mathcal{C}_1 , the intersections are well-behaved.

Lemma 2.7. *For all distinct $C_{ij}, C_{kl} \in \mathcal{C}_1$ we have*

$$|C_{ij} \cap C_{kl}| \leq d^2,$$

and for any two distinct points in \mathcal{P}_1 , there are at most d^2 curves in \mathcal{C} that contain both.

Proof. As just observed, we can assume that $i \neq j$ and $k \neq l$. The points $(q, q') \in C_{ij} \cap C_{kl}$ are on the intersection of the surface $C_2 \times C_2$ with the hyperplanes $H_{ij} : B(p_i, q) = B(p_j, q')$ and $H_{kl} : B(p_k, q) = B(p_l, q')$. Since, by definition of \mathcal{C}_1 , $|C_{ij} \cap C_{kl}|$ is finite, applying Bézout's inequality as in Lemma 2.6 shows that this intersection contains at most $\deg(C_2)^2 \cdot \deg(H_{ij}) \cdot \deg(H_{kl}) = d^2$ points.

The same argument gives $|\tilde{C}_{st} \cap \tilde{C}_{uv}| \leq d^2$ for all s, t, u, v with $(q_s, q_t) \neq (q_u, q_v) \in \mathcal{P}_1$. This is the dual statement to (q_s, q_t) and (q_u, q_v) lying in at most d^2 curves from \mathcal{C} . \square

Note that applying Bézout's inequality directly to these curves of degree at most d^2 gives $|C_{ij} \cap C_{kl}| \leq d^4$, which would lead to a worse degree dependence in our final bound.

Next we show that \mathcal{P}_0 and \mathcal{C}_0 are relatively small. We do this by showing that if two curves have infinite intersection, then this is related to an automorphism of C_2 , and by assumption C_2 does not have many automorphisms.

For a linear transformation $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$, we define its *graph on C_2* by

$$G_T = \{(q, q') \in \mathbb{C}^4 : q, q' \in C_2, T(q) = q'\}.$$

It is the intersection of the surface $C_2 \times C_2$ with the graph of T , which is a plane. Typically, these two surfaces in \mathbb{C}^4 would have finite intersection, but this is not always the case. When the intersection is infinite, this means that T is an automorphism of C_2 .

Lemma 2.8. *For any distinct $C_{ij}, C_{kl} \in \mathcal{C}$, there is a linear transformation T such that*

$$C_{ij} \cap C_{kl} = G_T.$$

If $|C_{ij} \cap C_{kl}| = \infty$, then T is an automorphism of C_2 , and we have $i \neq k$ and $j \neq l$.

The same statements hold for the dual curves \tilde{C}_{st} corresponding to points $(q_s, q_t) \in \mathcal{P}$.

Proof. If $(q, q') \in C_{ij} \cap C_{kl}$ then we have

$$\begin{aligned} B(p_i, q) &= B(p_j, q'), \\ B(p_k, q) &= B(p_l, q'), \end{aligned}$$

which we can rewrite as $N_{ik}q = N_{jl}q'$ with the matrices N_{ik}, N_{jl} from Definition 2.3. We have either $i \neq k$ or $j \neq l$; without loss of generality we assume $j \neq l$, so that N_{jl} is invertible by Lemma 2.4. We define a linear transformation T by

$$q' = T(q) = N_{jl}^{-1}N_{ik}q.$$

It follows that $C_{ij} \cap C_{kl} \subset G_T$. On the other hand, if $(q, q') \in G_T$, then $q, q' \in C_2$ and $q' = T(q) = N_{jl}^{-1}N_{ik}q$, so $N_{jl}q' = N_{ik}q$. This exactly means that $(q, q') \in C_{ij} \cap C_{kl}$, so in fact we have $C_{ij} \cap C_{kl} = G_T$. This proves the first statement of the lemma.

If $|C_{ij} \cap C_{kl}| = \infty$, then $|T(C_2) \cap C_2| = \infty$. Since C_2 and $T(C_2)$ are irreducible algebraic curves, Bézout's inequality implies that $T(C_2) = C_2$, i.e., T is an automorphism of C_2 .

Suppose $i = k$. If there are infinitely many points $(q, q') \in C_{ij} \cap C_{il}$, then they satisfy $N_{ii}q = N_{jl}q'$. Since N_{ii} is singular and its image is the line $y = x$, the same must be true for N_{jl} , which implies that $j = l$. Similarly, if $j = l$ and $|C_{ij} \cap C_{kl}| = \infty$, we get $i = k$.

The same arguments give the corresponding statements for the dual curves. \square

2.5 Incidence bound

To get an upper bound for the incidences between \mathcal{P}_1 and \mathcal{C}_1 , we use the following theorem, which we deduce from a theorem proved by Solymosi and De Zeeuw in [11].

Theorem 2.9. *Let $A, B \subset \mathbb{C}^2$ with $|A| = |B| = \mu$, let $\Pi \subset A \times B$, and let Γ be a set of algebraic curves in \mathbb{C}^4 of degree at most δ , with $|\Gamma| = \mu^2$. If any two points of Π are contained in at most Δ curves of Γ , then we have*

$$I(\Pi, \Gamma) = O\left(\delta^{4/3} \Delta^{1/3} \mu^{8/3}\right).$$

Proof. Theorem 1 and Remark 15 from [11] give this statement for curves in \mathbb{C}^2 . We can reduce to that case using a generic projection argument, for instance as worked out in detail in [8]. We will only sketch how that argument can be adapted to this situation.

Let $\psi : \mathbb{C}^4 \rightarrow \mathbb{C}^2$ be the projection $(z_1, z_2, z_3, z_4) \mapsto (z_1, z_3)$. We claim that there is a linear transformation $\varphi : \mathbb{C}^4 \rightarrow \mathbb{C}^4$ with a matrix of the form

$$\begin{pmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a' & b' \\ 0 & 0 & c' & d' \end{pmatrix},$$

so that $\pi := \psi \circ \varphi$ has the following properties: π is bijective on Π ; π induces a bijection between $I(\Pi, \Gamma)$ and $I(\pi(\Pi), \pi(\Gamma))$; for $\gamma, \gamma' \in \Gamma$, $\pi(\gamma)$ and $\pi(\gamma')$ are distinct algebraic curves in \mathbb{C}^2 . Because of the form of the matrix, we can write $\pi(\Pi) = A' \times B'$ with two sets $A', B' \subset \mathbb{C}$. The linear map does not increase the degree of the curves. Applying the main theorem of [11] gives the desired bound.

The claim is proved exactly as in [8, Corollary 2.5], by showing that the set of φ for which one of these properties fails is a lower-dimensional subset of the 8-dimensional space of such matrices. \square

By Lemma 2.7, \mathcal{P}_1 and \mathcal{C}_1 almost exactly satisfy the conditions of Theorem 2.9 with $A = B = T^*$, $\mu = m$, $\delta = d^2$, and $\Delta = d^2$; only the condition $|\mathcal{C}_1| = m^2$ need not quite hold, but it is easily forced by adding in dummy curves or points, without adding incidences. Thus we get the following bound.

Lemma 2.10. *We have the incidence bound*

$$I(\mathcal{P}_1, \mathcal{C}_1) = O\left(d^{10/3} m^{8/3}\right).$$

2.6 Conclusion

We show that the incidences coming from \mathcal{P}_0 and \mathcal{C}_0 are negligible.

Lemma 2.11. *If each of C_1, C_2 has $O(d^2)$ automorphisms, then*

$$|\mathcal{C}_0| = O(d^2 m) \quad \text{and} \quad |\mathcal{P}_0| = O(d^2 m).$$

Proof. We define a graph with vertices $C_{ij} \in \mathcal{C}_0$ and an edge between C_{ij} and C_{kl} if and only if $|C_{ij} \cap C_{kl}| = \infty$. We color an edge $C_{ij}C_{kl}$ with the transformation T if $C_{ij} \cap C_{kl} = G_T$; by Lemma 2.8, there is such a T for every edge.

If two edges of the form $C_{ij}C_{kl}$ and $C_{ij'}C_{k'l'}$ have the same color T , then $C_{ij} \cap C_{kl} = G_T = C_{ij'} \cap C_{k'l'}$. Then $G_T \subset C_{ij} \cap C_{ij'}$, so $|C_{ij} \cap C_{ij'}| = \infty$, contradicting Lemma 2.8.¹

It follows that every color T occurs at most m times, since for each i there is at most one j such that C_{ij} is incident with an edge of color T . By assumption, C has $O(d^2)$ automorphisms, so there are at most $O(d^2)$ colors, hence the graph has $O(d^2m)$ edges. By definition of \mathcal{C}_0 there are no isolated vertices, so the number of vertices is at most twice the number of edges, hence $|\mathcal{C}_0| = O(d^2m)$.

A similar argument applied to the dual curves gives the bound on $|\mathcal{P}_0|$. \square

Lemma 2.12. *If each of C_1, C_2 has $O(d^2)$ automorphisms, then*

$$I(\mathcal{P}, \mathcal{C}_0) = O(d^3m^2) \quad \text{and} \quad I(\mathcal{P}_0, \mathcal{C}) = O(d^3m^2).$$

Proof. Any C_{ij} has at most dm incidences with points $(q_s, q_t) \in \mathcal{P}$. This is because for any of the m choices for q_s , the corresponding q_t must be an intersection point of C_2 with the line $\{q \in \mathbb{C}^2 : B(p_j, q) = B(p_i, q_s)\}$. Since we assumed that C_2 is not a line, by Bézout's inequality there are at most d such intersection points.

Since $|\mathcal{C}_0| = O(d^2m)$, this gives $I(\mathcal{P}, \mathcal{C}_0) = O(d^3m^2)$. The dual argument gives the second bound. \square

We get the overall incidence bound

$$I(\mathcal{P}, \mathcal{C}) \leq I(\mathcal{P}_0, \mathcal{C}) + I(\mathcal{P}, \mathcal{C}_0) + I(\mathcal{P}_1, \mathcal{C}_1) = O(d^{10/3}m^{8/3}).$$

Combining this with $I(\mathcal{P}, \mathcal{C}) = |\mathcal{Q}| \geq m^4/|\mathcal{B}|$ from Lemma 2.5 and $m \geq n/d$ gives

$$|\mathcal{B}(S, T)| = \Omega(|\mathcal{B}|) = \Omega(m^4/|\mathcal{Q}|) = \Omega(d^{-10/3}m^{4/3}) = \Omega(d^{-14/3}n^{4/3}),$$

which completes the proof of Theorem 2.1.

¹In fact, the edges of the same color form a clique, but we do not need this fact.

3 Linear automorphisms

In this section we study algebraic curves that have infinitely many linear automorphisms. Although the topic seems classical, we were not able to find in the literature the exact statement that we need, so we provide our own proof.

Recall that by a *(linear) automorphism* of a curve C we mean an invertible linear transformation $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ such that $T(C) = C$. Note that in algebraic geometry, “automorphism” often denotes a *polynomial* transformation (or “morphism”) that fixes the curve, or sometimes a *projective* transformation that fixes the curve. The classic theorem about polynomial automorphisms is Hurwitz’s Theorem, which states that a nonsingular curve of genus $g \geq 2$ has at most $84(g - 1)$ polynomial automorphisms (see for instance [4], Exercise IV.2.5). If C has degree d , then we have $g \leq d^2$, so we get a bound in terms of the degree d . However, this does not give the exact picture for linear automorphisms. For nonsingular curves, it would reduce the question to conics, for which one can easily compute what the linear automorphisms are. However, there are many higher-degree singular curves of genus 0, for which it is harder to determine the linear automorphisms. This is what we do directly with an elementary approach, sidestepping Hurwitz’s Theorem (and its difficult proof) altogether.

The theorem we prove in this section is the following. Together with Corollary 2.2, it implies Theorem 1.3. Special curves are defined in Definition 1.1.

Theorem 3.1. *An irreducible algebraic curve of degree d has $O(d^2)$ linear automorphisms, unless it is a special curve.*

Proof. The theorem follows directly from Lemma 3.3 and Lemma 3.4 below. Assume C is not a special curve. If C does not contain the origin, it has $O(d^2)$ automorphisms by Lemma 3.3, and if it does contain the origin, it has $O(d^2)$ automorphisms by Lemma 3.4. \square

Example 3.2. *Special curves have infinitely many automorphisms. For $x^k = y^\ell$, the matrix*

$$\begin{pmatrix} \alpha^\ell & 0 \\ 0 & \alpha^k \end{pmatrix}$$

defines an automorphism for all $\alpha \in \mathbb{C} \setminus \{0\}$. It then clearly follows that a linearly equivalent curve has infinitely many automorphisms.

An initial idea for proving Theorem 3.1 would be to observe the following about an automorphism T of the curve C . If L is an eigenline of T and $q \in C \cap L$, then $T^i(q) \in C \cap L$ for all i . If the eigenvalue of L is not a root of unity, then the points $T^i(q)$ would form an infinite set in $C \cap L$, so by Bézout’s inequality, C would have to equal L . However, this approach fails, because $C \cap L$ may be empty (and this is indeed what happens for special curves). We therefore have to use a similar but trickier argument. Over \mathbb{R} , the argument would be considerably simpler, as we would not have to worry about roots of unity.

Our proof of Theorem 3.1 rests on the three lemmas below. The first two are complementary and together imply Theorem 3.1. The third, more technical, lemma is used in the proofs of the first two lemmas to handle specific subcases. We use some concepts from the theory of algebraic curves, for which we refer to [4]; namely the projective plane, singularities and their branches, and intersection multiplicity.

In these lemmas we let C be an irreducible algebraic curve of degree d , and f a minimum-degree polynomial with $C = Z(f)$. We write T_λ for the scaling transformation defined by $T_\lambda(p) = \lambda p$, with $\lambda \in \mathbb{C} \setminus \{0\}$. We write L_m for the line $y = mx$ with $m \in \mathbb{C}$.

Lemma 3.3. *Suppose C is not a line and does not contain the origin. Then C has $O(d^2)$ automorphisms, unless it is linearly equivalent to $x^k y^\ell = 1$, with $k, \ell \geq 1$.*

Proof. Suppose C has more than d^2 automorphisms, and choose matrices A_1, \dots, A_{d^2+1} from among them.² We claim that for all but finitely many $m \in \mathbb{C}$, the line L_m has the following two properties: $|C \cap L_m| = d$, and the lines $A_i L_m$ are distinct. The first property fails only for the finitely many m such that L_m is tangent to C , or intersects C at infinity or in a singularity. The second property fails only when for some pair i, j , L_m is a line such that $(A_i - A_j)L_m = 0$; if such a line exists, it is unique.

Choose L_m with the two properties above. Suppose $q, \lambda q \in C \cap L_m$ for some $\lambda \in \mathbb{C} \setminus \{0, 1\}$. Then the points $A_i q, \lambda A_i q$ are all on C , and they are all distinct by the second property of L_m . Since T_λ sends $A_i q$ to $\lambda A_i q$, the irreducible curves C and $T_\lambda(C)$ have $d^2 + 1$ points in common, so by Bézout's inequality we have $T_\lambda(C) = C$. Thus T_λ is an automorphism of C , and $T_\lambda^i q = \lambda^i q$ lies on $L_m \cap C$ for all $i \in \mathbb{Z}$. If more than d of the numbers λ^i are distinct, then Bézout's inequality gives $C = L_m$. Otherwise, λ is a root of unity of order at most d .

Choose $q \in C \cap L_m$ and consider the argument in the previous paragraph for q together with each of the $d - 1$ other points in $C \cap L_m$ in the role of λq . This, together with $\lambda = 1$, gives d distinct values of λ , each of which is a root of unity of order at most d . This implies that one of these λ is a primitive d -th root of unity, i.e., $\lambda^d = 1$ but $\lambda^k \neq 1$ for $0 < k < d$.

Let T_λ an automorphism of C with λ a primitive d -th root of unity. Write $q = (q_x, q_y)$. Then, for any m as above, $\lambda^i q_x$ must be a root of $f(x, mx)$ for each $i = 0, \dots, d - 1$. Thus

$$f(x, mx) = \alpha \prod_{i=0}^{d-1} (x - \lambda^i q_x) = \alpha (x^d - q_x^d)$$

for some $\alpha \in \mathbb{C} \setminus \{0\}$. Because this holds for all but finitely many m , it follows that $f(x, y)$ only has terms of degree 0 or d , and (after scaling) there are $a_i, c \in \mathbb{C}$ such that

$$f(x, y) = \prod_{i=1}^d (y - a_i x) + c.$$

The lines L_{a_i} are the asymptotes of C , and any automorphism of C must permute these lines (i.e., it must permute the set $\{L_{a_i}\}$). In Lemma 3.5 we will show that C has $O(d^2)$ automorphisms if it permutes a set of three or more lines, so we are done if at least three a_i are distinct. Otherwise, only two of the a_i are distinct, which means that

$$f(x, y) = (y - b_1 x)^k (y - b_2 x)^\ell + c$$

for some integers k, ℓ and $b_1, b_2, c \in \mathbb{C}$. This equation is linearly equivalent to $x^k y^\ell = 1$. \square

Lemma 3.4. *Suppose C is not a line and contains the origin. Then C has $O(d^2)$ automorphisms, unless it is linearly equivalent to $x^k = y^\ell$, with $k, \ell \geq 1$.*

Proof. Now C need not have exactly d distinct points on most lines L_m , since if it has a singularity at the origin, it may have high intersection multiplicity with all lines at the origin. However, there is a $k \leq d$ such that most lines have $|L_m \cap C| = k$. By the same argument

²We really mean d^2 ; the $O(d^2)$ in the lemma comes from the second part of this proof

as in Lemma 3.3, we can reduce to the case where T_λ is an automorphism, with $\lambda^k = 1$, and $C \cap L_m$ consists of the points $\lambda^i q$ for $i = 0, \dots, k$. Hence, for most L_m we have

$$f(x, mx) = (\alpha x^k + \beta)x^{d-k} = \alpha x^d + \beta x^{d-k},$$

and it follows that

$$f(x, y) = a \prod_{i=1}^d (y - a_i x) + b \prod_{j=1}^{d-k} (y - b_j x).$$

Any automorphism must permute the asymptotes L_{a_i} , and it must also permute the lines L_{b_j} , because these are the tangent lines of C at the origin. Note that the lines L_{b_j} are distinct from the lines L_{a_i} because f is irreducible. By Lemma 3.5, if at least three of all these lines together are distinct, then C has $O(d^2)$ automorphisms. Otherwise, we must have all a_i equal and all b_j equal, so

$$f(x, y) = a(y - a'x)^d + b(y - b'x)^{d-k},$$

which is linearly equivalent to $x^d = y^{d-k}$. \square

Lemma 3.5. *Let \mathcal{L} be a set of lines through the origin in \mathbb{C}^2 , with $3 \leq |\mathcal{L}| \leq 2d$. Then an algebraic curve $C \subset \mathbb{C}^2$ has $O(d^2)$ automorphisms that permute \mathcal{L} .*

Proof. We work in the projective plane. Let L_∞ be the line at infinity and P_∞ the set of points at infinity of the lines in \mathcal{L} , so $3 \leq |P_\infty| \leq 2d$. For a linear T on \mathbb{C}^2 we write φ_T for the Möbius transformation that T induces on L_∞ . We note that any such Möbius transformation is determined by its image on any three points. Let G be the group of automorphisms of C that permute \mathcal{L} , and $G_\infty := \{\varphi_T : T \in G\}$, so every $\varphi \in G_\infty$ permutes P_∞ .

We first note that G and G_∞ are finite groups. Since $|P_\infty| \geq 3$, a permutation of P_∞ corresponds to at most one transformation in G_∞ , which implies that G_∞ is finite. To show that G is finite, we show that for any $\varphi \in G_\infty$ there are finitely many $T \in G$ such that $\varphi_T = \varphi$. Choose two points of C in \mathbb{C}^2 that do not lie on the same line through the origin. Then for a fixed $\varphi \in G_\infty$, any $T \in G$ with $\varphi_T = \varphi$ must send these two points to points on two fixed lines, and given the images of these two points, T is determined. Since C has at most d points on these lines, there are finitely many possible images for these points, which implies that there are at most finitely many such T . Thus G is also finite.³

We now use some basic facts about Möbius transformations, which can be found in for instance [7, Chapter 3] or [6, Chapter 2]. A Möbius transformation of finite order has exactly two fixed points. A finite subgroup of the group of Möbius transformations is either a cyclic group, a dihedral group, or one of S_4 , A_4 , or A_5 (see [6, Corollary 2.13.7]), so G_∞ must be one of these groups. In the last three cases, G_∞ has size at most 60. If G_∞ is cyclic, then every $\varphi \in G_\infty$ has the same two fixed points. Since $|P_\infty| \geq 3$, we can choose a $p \in P_\infty$ that is not one of the two fixed points, and then choosing the image of p from the $|P_\infty| \leq 2d$ candidates determines φ . Thus $|G_\infty| \leq 2d$. If G_∞ is dihedral, there are two points such that any $\varphi \in G_\infty$ either fixes them, or swaps them. The same argument as for the cyclic case then gives that $|G_\infty| \leq 4d$. Altogether we have $|G_\infty| \leq \max\{4d, 60\} = O(d)$.

Fix $\varphi \in G_\infty$ and choose a point $q \in C$ on a line L through the origin that corresponds to a fixed point of φ . Then for any $T \in G$ with $\varphi_T = \varphi$, $T(q)$ must lie on a L , as well as on C . Since C is not a line, it has at most d points on L . Thus there are at most d choices for $T(q)$, and given this choice, T is determined. It follows that $|G| \leq d \cdot |G_\infty| = O(d^2)$. \square

³This rough argument already gives a bound on $|G|$, but it is too large for our purposes.

4 Discussion

Degree dependence. Let $F : \mathbb{C}^2 \times \mathbb{C}^2 \rightarrow \mathbb{C}$ be a polynomial function. Given a set S of n points in \mathbb{C}^2 , by interpolation there exists an algebraic curve of degree $O(n^{1/2})$ containing S . Thus, a bound $\Omega(d^{-\alpha}n^{1+\beta})$ for the number of distinct values of F on a curve gives a lower bound $|F(S)| = \Omega(n^{1+\beta-\alpha/2})$ on the number of distinct values of F .

In [8], where $F(p, q) = (p_x - q_x)^2 + (p_y - q_y)^2$ was the Euclidean distance function, the bound obtained (over \mathbb{R}) was $\Omega(d^{-11}n^{4/3})$, which clearly makes the interpolation argument above useless. Part of the goal for this paper was to see if this could be improved for bilinear forms. Over \mathbb{C} , our main bound from Theorem 1.3 also gives nothing. Over \mathbb{R} , our proof would give $\Omega(d^{-2}n^{4/3})$ (mainly because the dependence on d in the real equivalent of Theorem 2.9 would be better; see [11]). Then interpolation gives $|F(S)| = \Omega(n^{1/3})$, which is more tangible but still rather weak.

We conclude that to obtain an interesting bound from this interpolation argument, one would have to improve the exponent $4/3$, or the dependence on d in Theorem 2.9.

Elekes-Rónyai on curves. Our result fits into the general framework of Elekes and Rónyai [2], which considers polynomial functions

$$F : X_1 \times X_2 \rightarrow X_3,$$

for varieties X_1, X_2, X_3 of the same dimension. Elekes and Rónyai [2] consider the case where $X_1 = X_2 = X_3 = \mathbb{R}$, and proved that F takes $\omega(n)$ values, unless it has one of the special forms $F(x, y) = G(H(x) + K(y))$ or $F(x, y) = G(H(x) \cdot K(y))$ for polynomials G, H, K . The lower bound was improved by Raz, Sharir, and Solymosi [9] to $\Omega(n^{4/3})$.

In our case we have $X_1 = X_2 = C$ and $X_3 = \mathbb{C}$, and F a bilinear polynomial. We note that if M is not invertible, we have $B(p, q) = L_1(p) \cdot L_2(q)$ for linear polynomials L_1, L_2 , which one can see as an analog of the multiplicative form of Elekes and Rónyai (an additive form is actually not possible here). This (and other, unpublished, considerations) leads us to the following conjecture.

Conjecture 4.1. *Let $C \subset \mathbb{C}^2$ be an algebraic curve of degree d_C and $F : C \times C \rightarrow \mathbb{C}$ a polynomial of degree d_F . Then for any $S \subset C$ we have*

$$|F(S)| = \Omega_{d_C, d_F}(|S|^{4/3}),$$

unless $F(p, q) = G(H(p) + K(q))$, $F(p, q) = G(H(p) \cdot K(q))$, or unless C is rational.

It seems reasonable to take rational curves as exceptions in this statement, because these are the curves that can have infinitely many automorphisms defined by higher-degree polynomials (essentially by Hurwitz's Theorem, see Section 3). Of course, for specific functions the exact class of exceptions may be smaller.

When $F(p, q) = G(H(p) + K(q))$ or $F(p, q) = G(H(p) \cdot K(q))$, $|F(S_1, S_2)| = O(n)$ is possible for different sets S_1, S_2 . For the additive form, choose a set S_1 of intersection points of C with the curve $H(p) = i$ for $i = 1, \dots, |S|$, and a set S_2 of intersection points with $K(q) = j$ for $j = 1, \dots, n$ (this is certainly possible over \mathbb{C} ; over \mathbb{R} one needs to be more careful). For the multiplicative form, one can do the same with $H(p) = 2^i$ and $K(q) = 2^j$. However, it seems difficult to construct such an example with $S_1 = S_2$, unless $H = K$.

The exponent $4/3$. The exponent $4/3$ is not expected to be tight. In all of the papers [10, 8, 9] that obtain it in this framework, the main open problem is to improve this exponent, perhaps as far as $\Omega(|S|^{2-\varepsilon})$. In these proofs, the room for improvement seems to be in the incidence bound. Perhaps one can improve on the Szemerédi-Trotter-like exponent in Theorem 2.9 by using the specific nature of the incidence problem that one gets here, with the point set being a Cartesian product, and the curves being a very restricted family. Indeed, the curves are dual to a point set that is also a Cartesian product.

Acknowledgments.

Both authors were partially supported by Swiss National Science Foundation Grants 200020-144531 and 200021-137574. Part of this research was performed during the second author’s visit to the Institute for Pure and Applied Mathematics in Los Angeles, which is supported by the National Science Foundation. The authors thank János Pach for all his support.

References

- [1] M. Charalambides, *Distinct Distances on Curves Via Rigidity*, Discrete & Computational Geometry **51**, 666–701, 2014.
- [2] G. Elekes and L. Rónyai, *A combinatorial problem on polynomials and rational functions*, Journal of Combinatorial Theory, Series A **89**, 1–20, 2000.
- [3] J. Garibaldi, A. Iosevich, and S. Senger, *The Erdős Distance Problem*, AMS Student Library Series **56**, 2011.
- [4] R. Hartshorne, *Algebraic Geometry*, Springer-Verlag, 1977.
- [5] J. Heintz, *Definability and fast quantifier elimination in algebraically closed fields*, Theoretical Computer Science **24**, 239–277, 1983.
- [6] G. Jones and D. Singerman, *Complex functions*, Cambridge University Press, 1987.
- [7] T. Needham, *Visual Complex Analysis*, Oxford University Press, 1997.
- [8] J. Pach and F. de Zeeuw, *Distinct distances on algebraic curves in the plane*, Proceedings of the thirtieth annual symposium on Computational geometry, 549–557, 2014.
- [9] O.E. Raz, M. Sharir, and J. Solymosi, *Polynomials vanishing on grids: The Elekes-Rónyai problem revisited*, Proceedings of the thirtieth annual symposium on Computational geometry, 251–260, 2014.
- [10] M. Sharir, A. Sheffer, and J. Solymosi, *Distinct distances on two lines*, Journal of Combinatorial Theory, Series A **120**, 1732–1736, 2013.
- [11] J. Solymosi and F. de Zeeuw, *Incidence bounds for complex algebraic curves on Cartesian products*, arXiv:1502.05304, 2015.
- [12] S. Steinerberger, *A note on the number of different inner products generated by a finite set of vectors*, Discrete Mathematics **310**, 1112–1117, 2010.